

Piecewise Deterministic Markov Processes for Scalable Monte Carlo on Restricted Domains

Joris Bierkens¹, Alexandre Bouchard-Côté², Arnaud Doucet³, Andrew B. Duncan⁴, Paul Fearnhead⁵, Gareth Roberts⁶, and Sebastian J. Vollmer⁶

¹EEMCS, TU Delft, Netherlands.

²Department of Statistics, University of British Columbia, Canada.

³Department of Statistics, University of Oxford, U. K.

⁴School of Mathematical and Physical Sciences, University of Sussex, U. K.

⁵Department of Mathematics and Statistics, Lancaster University, U. K.

⁶Department of Statistics, University of Warwick, U. K.

January 17, 2017

Abstract

Piecewise deterministic Monte Carlo methods (PDMC) consist of a class of continuous-time Markov chain Monte Carlo methods (MCMC) which have recently been shown to hold considerable promise. Being non-reversible, the mixing properties of PDMC methods often significantly outperform classical reversible MCMC competitors. Moreover, in a Bayesian context they can use sub-sampling ideas, so that they need only access one data point per iteration, whilst still maintaining the true posterior distribution as their invariant distribution. However, current methods are limited to parameter spaces of the form \mathbb{R}^d . We show how these algorithms can be extended to applications involving restricted parameter spaces. In simulations we observe that the resulting algorithm is more efficient than Hamiltonian Monte Carlo for sampling from truncated logistic regression models. The theoretical framework used to justify this extension lays the foundation for the development of other novel PDMC algorithms.

1 Introduction

Markov chain Monte Carlo (MCMC) methods have been central to the wide-spread use of Bayesian methods. However their applicability to some modern applications has been limited due to their high computational cost, particularly in big-data, high-dimensional settings. This has led to interest in new MCMC methods, particularly non-reversible methods which can mix better than standard reversible MCMC [9], and

variants of MCMC that require accessing only small subsets of the data at each iteration [26].

One promising set of algorithms are based around simulating a piecewise deterministic Markov process [23, 7, 6]. Such processes explore the state space according to a constant velocity, where the velocity changes at random event times. The rate of these event times, and the change in velocity at each event, are chosen so that the resulting process has the posterior distribution as its invariant distribution. We will refer to this family of sampling methods as Piecewise Deterministic Monte Carlo methods (PDMC).

The PDMC methods proposed to date exhibit some remarkable properties of practical interest:

1. there is no rejection, i.e. each piece of computation implies exploration of new configurations for all variables of the model;
2. there is no need for symplectic integrators (which are required for Hamiltonian Monte Carlo [10]) because trajectories consist simply of straight lines;
3. the underlying Markov process is non-reversible;
4. factor graph decompositions of the target distribution can be leveraged to perform sparse updates of the variables [7, 19, 23];
5. and it is often possible to use only a small sub-sample of the data at each iteration, whilst still being guaranteed to target the true posterior distribution [5, 7, 13].

Existing PDMC algorithms can only be used to sample from posteriors where the parameters can take any value in \mathbb{R}^d . In this paper we show how to extend PDMC methodology to deal with constraints on the parameters. Such models are ubiquitous in machine learning and statistics. For example, many popular models used for binary, ordinal and polychotomous response data are multivariate real-valued latent variable models where the response is given by a deterministic function of the latent variables [1, 11, 24]; e.g. a binary response is set equal to 1 if a linear combination of the latent variables is positive and -1 otherwise. Under the posterior distribution, the domain of the latent variables is then constrained based on the values of the responses. Additional examples arise in regression where prior knowledge restricts the signs of marginal effects of explanatory variables such as in econometrics [14], image processing and spectral analysis [3], [15] and non-negative matrix factorization [17]. When performing Bayesian inference in these contexts, posterior simulation requires sampling from constrained distributions. A few methods for dealing with restricted domains are available but these either target an approximation of the correct distribution [22] or are limited in scope [21].

The general framework presented here is sufficiently rich to include both the *bouncy particle sampler* [7, 23], the *zig zag sampler* [5, 6] as well as *event chain* sampling methods of Werner Krauth and collaborators [4, 19], as we will illustrate with an example. We present simulation results of PDMC for isotonic logistic regression. In this case we see the PDMC methods can be substantially more efficient than Hamiltonian Monte Carlo.

Additionally, we also give a general presentation of PDMC methods, which shows how the choices of the dynamics of the process affect their stationary distribution. These results can aid with the development of novel PDMC algorithms.

2 Piecewise Deterministic Monte Carlo

Our objective is to compute expectations with respect to a probability distribution π on $\mathcal{O} \subseteq \mathbb{R}^d$ which is assumed to have smooth density, also denoted $\pi(x)$, with respect to the Lebesgue measure on \mathcal{O} .

We denote the state of our continuous-time Markov process at time t as $Z_t = (X_t, V_t)$, taking values in the domain $E = \mathcal{O} \times \mathcal{V}$, where \mathcal{O} and \mathcal{V} are subsets of \mathbb{R}^d , such that \mathcal{O} is open. The dynamics of Z_t are easy to describe if one views X_t as position and V_t as velocity. The position process X_t moves deterministically, with constant velocity V_t between a discrete set of *switching*

times which are simulated according to N inhomogeneous Poisson processes, with respective intensity functions $\lambda_i(x_t, v_t)$, $i = 1, \dots, N$, depending on the current state of the system. At each switching time the position stays the same, but the velocity is updated according to a specified transition kernel. More specifically, suppose the next switching event occurs from i^{th} Poisson process, then the velocity immediately after the switch is sampled randomly from the probability distribution $Q_i(x, v, \cdot)$, i.e. for any measurable subset $A \subset \mathcal{V}$, the probability that the new velocity is in A given the current position x and velocity v before the switch is $Q_i(x, v, A)$. The switching times are random, and designed in conjunction with the kernels $\{Q_i\}_{i=1, \dots, N}$ so that the invariant distribution of the process coincides with the target distribution π . An example of a realisation of this process is shown in Figure 1(f).

The resulting stochastic process is a Piecewise Deterministic Markov Process (PDMP, [8]). For it to be useful as the basis of a Piecewise Deterministic Monte Carlo (PDMC) algorithm we need to (i) be able to easily simulate this process; and (ii) have simple recipes for choosing the intensities, $\{\lambda_i\}_{i=1, \dots, N}$, and transition kernels, $\{Q_i\}_{i=1, \dots, N}$, such that the resulting process has $\pi(x)$ as its marginal stationary distribution. We will tackle each of these problems in turn.

2.1 Simulation

We shall first assume that $\mathcal{O} = \mathbb{R}^d$. The key challenge in simulating our PDMP is simulating the event times. The intensity of events is a function of the state of the process. But as the dynamics between event times are deterministic, we can easily represent the intensity for the next event as a deterministic function of time. Suppose that the PDMP is driven by a single inhomogeneous Poisson process with intensity function

$$\tilde{\lambda}(u; X_t, V_t) = \lambda(X_t + uV_t, V_t), \quad u \geq 0.$$

We can simulate the first event time directly if we have an explicit expression for the inverse function of the monotonically increasing function

$$u \mapsto \int_0^u \tilde{\lambda}(s; X_t, V_t) ds. \quad (1)$$

In this case the time until the next event is obtained by (i) simulating a realization, y say, of an exponential random variable with rate 1; and (ii) setting the time until the next event as the value τ that solves $\int_0^\tau \tilde{\lambda}(s; X_t, V_t) ds = y$.

In practice inverting Equation 1 is often not practical. In such cases simulation can be carried via *thinning* [18]. This requires finding a tractable upper

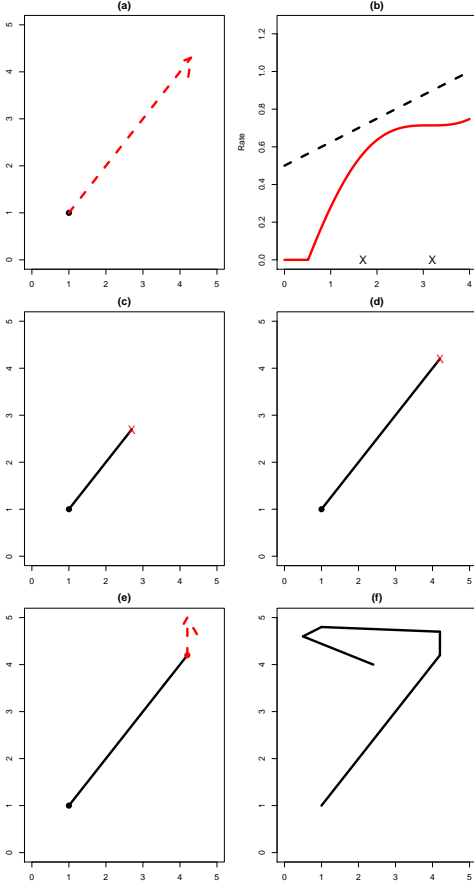


Figure 1: Schematics of a PDMC algorithm for a PDMP driven by a single inhomogeneous process with intensity $\tilde{\lambda}$, along with output for a two-dimensional target. (a) Initial position (shown by dot) and velocity (shown by arrow). (b) We first calculate $\tilde{\lambda}$ (red lower line), by evaluating the event rate along the current path of the process; and we bound this by a simple function $\bar{\lambda}$ (black upper line). We then propose event times from a Poisson process with rate $\bar{\lambda}$, the first two event times are shown by crosses. In practice these are simulated sequentially. (c) Given the first event time, u say, we simulate the deterministic process forward for a period of length u . We then accept or reject this event, with acceptance probability $\tilde{\lambda}(u)/\bar{\lambda}(u)$. Note that rejection in this context means the position is updated to move up to the rejected point, but the velocity is not changed. Here we reject, so (d) we then simulate the next proposed event-time, and repeat the accept-reject step. In this case we accept so (e) we then simulate a new velocity, and repeat the above iterations. (f) Example output of the process over a series of iterations.

bound on the rate, $\bar{\lambda}(u) \geq \tilde{\lambda}(u; X_t, V_t)$ for all $u > 0$. Such an upper bound will typically take the form of a piecewise linear function or a step function. Note that the upper bound $\bar{\lambda}$ is only required to be valid along the trajectory $u \mapsto (X_t + uV_t, V_t)$ in $\mathcal{O} \times \mathcal{V}$. Therefore the upper bound will depend on the starting point (X_t, V_t) of the line segment we are currently simulating.

We then propose potential events by simulating events from an inhomogeneous Poisson process with rate $\bar{\lambda}(u)$, and accept an event at time u with probability $\tilde{\lambda}(u; X_t, V_t)/\bar{\lambda}(u)$. The time of the first accepted event will be the time until the next event in our PDMP.

For a PDMP driven by N inhomogeneous Poisson processes with intensities $\{\lambda_i\}_{i=1}^N$ the previous steps lead to the following algorithm for simulating the next event of our PDMP. This algorithm can be iterated to simulate the PDMP for a chosen number of events or a pre-specified time-interval.

- (0) **Initialize:** Set t to the current time and (X_t, V_t) to the current position and velocity.
- (1) **Determine bound:** For each $i \in 1, \dots, N$, find a convenient function $\bar{\lambda}_i$ satisfying $\bar{\lambda}_i(u) \geq \tilde{\lambda}_i(u; X_t, V_t)$ for all $u \geq 0$, depending on the initial point (X_t, V_t) from which we are departing.
- (2) **Propose event:** For $i = 1, \dots, N$, simulate the first event times τ'_i of a Poisson process with rate function $\bar{\lambda}_i$.
- (3) Let $i_{\min} = \arg \min_{j=1, \dots, N} \tau'_j$ and $\tau' = \tau'_{i_{\min}}$.
- (4) **Accept/Reject event:** With probability

$$\frac{\tilde{\lambda}_i(\tau'; X_t, V_t)}{\bar{\lambda}_i(\tau')}$$

accept this event.

- (4.1) **Upon acceptance:** set $X_{t+\tau'} = X_t + \tau'V_t$; simulate a new velocity V' from $Q_{i_{\min}}(X_{t+\tau'}, V_t, \cdot)$ and set $V_{t+\tau'} = V'$.
- (4.2) **Upon rejection:** set $X_{t+\tau'} = X_t + \tau'V_t$ and set $V_{t+\tau'} = V_t$.
- (5) **Update:** Record the time $t + \tau'$ and state $(X_{t+\tau'}, V_{t+\tau'})$.

A schematic of this algorithm is shown in Figure 1.

2.2 Output of PDMC Algorithms

The output of these algorithms will be a sequence of event times $t_1, t_2, t_3, \dots, t_K$ and associated states $(X_1, V_1), (X_2, V_2), \dots, (X_K, V_K)$. To obtain the value

of the process at times $t \in [t_k, t_{k+1})$, we can linearly interpolate the continuous path of the process between event times, i.e. $X_t = X_{t_k} + V_k(t - t_k)$. Time integrals $\int_0^t f(X_s) ds$ of a function f of the process X_t can be approximated by numerically integrating the one dimensional integral along the piecewise linear trajectory of the PDMP. However, in many instances, such integrals can be computed analytically from the output of the above algorithm. For example, for $f(x) = x^{(r)}$, the r^{th} component of x we obtain,

$$\int_0^{t_K} X_s^{(r)} ds = \sum_{k=1}^K \left(X_{t_{k-1}}^{(r)} \tau_k + V_{t_{k-1}}^{(r)} \frac{\tau_k^2}{2} \right),$$

where $\tau_k = t_k - t_{k-1}$. Alternatively we can sample the PDMP at a set of evenly spaced time points along the trajectory and use this collection as an approximate sample from our target distribution. Note that using the values of the process at event times will not produce samples drawn from the stationary distribution of the PDMP [8].

It can be shown from the above construction that Z_t is a strong Markov process. Additional conditions are required to ensure that Z_t is ‘non-explosive’, in the sense that, almost surely, there are only finitely many switches in every finite time interval. The result below provides on such condition, valid for compact \mathcal{V} .

Proposition 1. *Suppose that \mathcal{V} is a compact subset of \mathbb{R}^d and that the switching intensities $\{\lambda_i\}_{i=1}^N$ are piecewise continuous, then the process $Z_t = (X_t, V_t)$ defined above is a non-explosive process.*

The proof can be found in the supplementary material.

2.3 Choosing the Intensity and Transition Kernel

Most existing PDMC methods [5, 7, 23, 7] assume that the target density, $\pi(x)$, is defined on \mathbb{R}^d . They further require that $\log \pi(x)$ be differentiable. Assuming these two conditions hold, we can provide criteria which must hold for a given probability distribution to be a stationary distribution of Z_t . We shall consider stationary distributions for which x and v are independent, i.e. distributions of the form $\pi(x)dx \otimes \rho(dv)$ on E . Furthermore we assume that

$$\pi(x) \propto \exp(-U(x)) \quad (2)$$

where U is continuously differentiable.

We impose the following condition:

$$\begin{aligned} & \int_{\mathcal{V}} \int_U \sum_{i=1}^N \lambda_i(x, v) Q_i(x, v, du) \rho(dv) \\ &= \int_U \int_{\mathcal{V}} \sum_{i=1}^N \lambda_i(x, v) Q_i(x, u, dv) \rho(dv), \end{aligned} \quad (3)$$

for all $x \in \mathcal{O}$ and $U \subset \mathcal{V}$ measurable. This is implied by the easier but stronger condition that each Q_i is reversible with respect to ρ , i.e. for $i = 1, \dots, N$,

$$Q_i(x, u, dv) \rho(dv) = Q_i(x, v, du) \rho(du). \quad (4)$$

Moreover, we shall require the following condition which relates the probability flow with the switching intensities λ_i :

$$\begin{aligned} & \sum_{i=1}^N \int_{\mathcal{V}} \lambda_i(x, v) Q_i(x, u, dv) - \sum_{i=1}^N \lambda_i(x, u) \\ &= -u \cdot \nabla U(x), \end{aligned} \quad (5)$$

for all $(x, u) \in E$. The following result is proved in the supplementary material.

Proposition 2. *Suppose that conditions (3) and (5) hold. Then the probability distribution $\pi(x) dx \otimes \rho(dv)$ is an invariant distribution of the process Z_t .*

Under the assumption that the resulting PDMP is ergodic (for sufficient conditions see e.g. [5, 7]), we have the following version of the law of large numbers [16] for the PDMP $Z_t = (X_t, V_t)$: For all square integrable functions $f : \mathbb{R}^d \rightarrow \mathbb{R}$ (i.e. $\int_{\mathbb{R}^d} f^2(x) \pi(x) dx < \infty$), we have that, with probability one,

$$\int_{\mathbb{R}^d} f(x) \pi(x) dx = \lim_{T \rightarrow \infty} \frac{1}{T} \int_0^T f(X_s) ds.$$

It is this formula which allows us to use PDMPs for Monte Carlo purposes.

2.4 Examples

Current PDMC algorithms differ in terms of how the Q_i and λ_i are chosen such that the above equation holds for some simple distribution for the velocity.

In the following examples δ_x denotes the Dirac-measure centered in x .

The Zig-Zag sampler [5] can be recovered by choosing $N = d$ and picking as velocity space $\mathcal{V} = \{-1, +1\}^d$ equipped with discrete uniform distribution ρ , defining switching rates $\lambda_i(x, v) = \max(v_i \partial_{x_i} U(x), 0)$. The corresponding switching kernels over new directions are given by

$$Q_i(x, v, dv') = \delta_{F_i v}(dv'),$$

where $F_i : \mathcal{V} \rightarrow \mathcal{V}$ denotes the operation of flipping the i -th component, i.e. $(F_i v)(i) = -v(i)$, and $(F_i v)(j) = v(j)$ for $j \neq i$.

The Bouncy Particle Sampler introduced in [23] and explored in [7] is obtained setting $N = 1$ and $\rho = \mathcal{N}(0, I)$ on \mathbb{R}^d or $\rho = \mathcal{U}(S^{d-1})$ the uniform distribution on the unit sphere. The single switching rate is chosen to be $\lambda_{\text{BPS}}(x, v) = \max(v \cdot \nabla U(x), 0)$, with corresponding switching kernel Q which reflects v with respect to the orthogonal complement of ∇U with probability 1:

$$Q(x, v, dv') = \delta_{(I-2P_{\nabla U})v}(dv'),$$

where $P_y : z \mapsto \frac{z \cdot y}{\|y\|^2} y$ denotes orthogonal projection along the one dimensional subspace spanned by y .

As noted in [7] this algorithm suffers from reducibility issues. These can be overcome by refreshing the velocity by drawing a new velocity independently from $\rho(dv)$. In the simplest case the refreshment times come from an independent Poisson process with constant rate λ_{ref} . This also fits in the framework above by choosing $\lambda = \lambda_{\text{BPS}} + \lambda_{\text{ref}}$ and

$$Q(x, u, dv) = \frac{\lambda_{\text{BPS}}}{\lambda_{\text{BPS}} + \lambda_{\text{ref}}} \delta_{(I-2P_{\nabla U})u}(dv) + \frac{\lambda_{\text{ref}}}{\lambda_{\text{BPS}} + \lambda_{\text{ref}}} \rho(dv).$$

As a generalization of the BPS, one can consider a *preconditioned* version, by introducing a constant positive definite symmetric matrix M to rescale the velocity process, allowing certain directions to be preferred. More specifically, we set $\mathcal{V} = \mathbb{R}^d$, $\rho \equiv \mathcal{N}(0, M)$ and

$$Q(x, v, dv') = \delta_{(I-2P_{\nabla U}^M)v}(dv'),$$

where $P_y^M : z \mapsto \frac{z \cdot y}{y \cdot M y} M y$. It is straightforward to check that conditions (3) and (5) hold, and so $\pi \otimes \rho$ is an invariant measure. The choice of M plays a very similar role to the mass matrix in HMC, and careful tuning can give rise to dramatic increases in performance [13]. For Gaussian targets, the natural choice of M is the covariance of the target distribution. See [20] for a related approach.

3 Sampling over Restricted Domains

Suppose now that \mathcal{O} is an open pathwise-connected subset of \mathbb{R}^d with Lipschitz boundary $\partial\mathcal{O}$. We can adapt the PDMP scheme described previously to sample from a distribution $\pi(x)$ supported on \mathcal{O} . To ensure that x_t remains confined within \mathcal{O} the velocity of the process is updated whenever x_t hits $\partial\mathcal{O}$ so that the

process moves back into \mathcal{O} . We shall refer to such updates as *boundary reflections* even though they need not be specular reflections.

To handle boundary reflections, at every given time t , we keep track of the next reflection event in the absence of a switching event, i.e. we compute

$$\tau_b = \inf \{u > 0 : X_t + uV_t \notin \mathcal{O}\}.$$

If the boundary $\partial\mathcal{O}$ can be represented as a finite set of M hyper-planes in \mathbb{R}^d , then the cost of computing τ_b is $O(Md)$. When generating the switching event times and positions for Z_t we determine whether a boundary reflection will occur before the next potential switching event. If so, then we induce a switching event at time $t + \tau_b$ where $X_{t+\tau_b} \in \partial\mathcal{O}$ and sample a new velocity from the transition kernel Q_b , i.e. $V_{t+\tau_b} \sim Q_b(X_{t+\tau_b}, V_t, \cdot)$.

Although theoretically we may choose a new velocity pointing outwards and have an immediate second jump, we will for algorithmic purposes assume that the probability measure $Q_b(x, u, \cdot)$ for $(x, u) \in \partial\mathcal{O} \times \mathcal{V}$ is concentrated on those directions v for which $(v \cdot n(x)) \leq 0$, where $n(x)$ is the outward normal at $x \in \partial\mathcal{O}$.

Based on this, we thus modify steps (2)–(4) of the PDMP scheme as follows.

(2') **Propose event:** For $i = 1, \dots, N$ simulate the first event times τ'_i of a Poisson process with rate function $\bar{\lambda}_i$. Compute the next boundary reflection time τ_b .

(3') Let $i_{\min} = \arg \min_{j=1, \dots, N} \tau'_j$ and $\tau' = \tau'_{i_{\min}}$.

(4') **Accept/Reject event:**

(4.1') If $\tau_b < \tau'$ then set $\tau = \tau_b$; set $X_{t+\tau} = X_t + \tau V_t$; sample a new velocity $V_{t+\tau} \sim Q_b(X_{t+\tau}, V_t, \cdot)$.

(4.2') Otherwise with probability

$$\frac{\tilde{\lambda}_{i_{\min}}(\tau'; X_t, V_t)}{\bar{\lambda}_{i_{\min}}(\tau')}$$

accept the event at time $\tau = \tau'$.

(4.1.1') **Upon acceptance:** set $X_{t+\tau} = X_t + \tau V_t$; sample a new velocity $V_{t+\tau} \sim Q_{i_{\min}}(X_{t+\tau}, V_t, \cdot)$.

(4.2.2') **Upon rejection:** set $X_{t+\tau} = X_t + \tau V_t$ and set $V_{t+\tau} = V_t$.

To ensure that the process $\pi(x) dx \otimes \rho(dv)$ is an invariant distribution of Z_t we must impose conditions on the boundary transition kernel Q_b . The following result provides one such sufficient condition. The proof is deferred to the supplementary material.

Proposition 3. Consider the process Z_t on $\mathcal{O} \times \mathcal{V}$ where \mathcal{O} is an open, pathwise connected subset of \mathbb{R}^d with Lipschitz boundary. For $x \in \partial\mathcal{O}$, denote by $n(x)$ the outward unit normal of $\partial\mathcal{O}$. Suppose that conditions (3) and (5) hold on \mathcal{O} , and for $x \in \partial\mathcal{O}$,

$$Q_b(x, u, dv)\rho(du) = Q_b(x, v, du)\rho(dv) \quad (6)$$

and moreover

$$\int_{\mathcal{V}} (n(x) \cdot u) Q_b(x, v, du) = -v \cdot n(x) \quad (7)$$

for all $(x, v) \in \partial\mathcal{O} \times \mathcal{V}$. Then $\pi(x) dx \otimes \rho(dv)$ is an invariant distribution for the process Z_t .

In practice we only have to consider the exit region $\Gamma \subset \mathcal{O} \times \mathcal{V}$. For example if $\mathcal{O} = (a, b) \subset \mathbb{R}^1$ and $\mathcal{V} = \{-1, +1\}$, then $\Gamma = \{b, +1\} \cup \{a, -1\}$. The specification of Q_b on $(\partial\mathcal{O} \times \mathcal{V}) \setminus \Gamma$ is irrelevant as these points are never reached by Z_t . On this irrelevant set, we may choose Q_b as desired to satisfy (7). In practice it therefore suffices to check (4) and (7) on Γ .

Another possible behavior at the boundary is to generate the new reflected direction independently of the angle of incidence. This will also preserve the invariant distribution provided that ρ is isotropic.

Proposition 4. Consider the process Z_t as in the previous proposition, such that conditions (5) and (3) hold and the distribution ρ has mean zero. Then $\pi(x) dx \otimes \rho(dv)$ will be an invariant distribution for the process Z_t if $Q_b(x, v, du)$ is independent of v for all $x \in \partial\mathcal{O}$.

3.1 Example: Bouncy Particle Sampler

For the Bouncy Particle Sampler [23, 7], it is natural to choose

$$Q_b(s, v, du) = \delta_{(I - 2P_{n(s)})v}(du),$$

for $s \in \partial\mathcal{O}$, so that the process X_t reflects specularly at the boundary (i.e. angle of incidence equals angle of reflection of process with respect to the boundary normal). It is straightforward to check that condition (3) holds at the boundary and that (7) is satisfied. Similarly, for the preconditioned BPS with mass matrix, the natural choice of reflection at the boundary is given by

$$Q_b(s, v, du) = \delta_{(I - 2P_{n(s)})v}^M(du),$$

for which it is straightforward to confirm that the conditions of Proposition 3 are satisfied.

3.2 Example: Event Chains of Krauth et al.

Krauth and collaborators [4, 19] introduced the notion of event chains for simulations in statistical mechanics, with superior performance over other methods. We show in a simple example how this can be incorporated into the current framework.

Consider the simulation of two one-dimensional ‘hard disks’ of radius R on \mathbb{T}^1 , the one dimensional torus, i.e. the interval $[0, 1]$ with the sides identified. The locations of the centers are denoted by x_1 and x_2 . We have that $x = (x_1, x_2) \in \mathcal{O} := \{(x_1, x_2) \in \mathbb{T}^2 : d(x_1, x_2) > 2R\}$, where d represents the distance function on \mathbb{T}^1 . The target distribution is the uniform distribution on \mathcal{O} .

The *Straight Event Chain (SEC)* algorithm [4] initially samples a fixed direction vector $v \in \{e_1, -e_1, e_2, -e_2\}$ (with e_1, e_2 denoting canonical basis vectors) and then moves x along the direction v . This corresponds to one moving disk and one stationary disk. When the moving disk hits the stationary disk, its motion is transferred to the other disk in the same direction.

Since there are no changes of direction unless the disks hit each other, we have $\lambda \equiv 0$.




Figure 2: Illustration of \mathcal{O} , the point x at the boundary and the outward normal $n(x)$ for the event chain discussed in Section 3.2. The blue area is the exterior of \mathcal{O} as a subset of \mathbb{T}^2 . The width of the blue strip (in horizontal or vertical direction) is equal to $2R$.

Consider for simplicity a single event where we hit the boundary $\partial\mathcal{O}$ in such a way that $x_1 = x_2 + 2R \bmod 1$ (the other possibility being $x_2 = x_1 + 2R \bmod 1$). See Figure 2 for an illustration. At this

point x , the dynamics described above result in

$$Q_b(x, u, dv) = \begin{cases} \delta_{-e_2}(dv) & \text{if } u = -e_1, \\ \delta_{e_1}(dv) & \text{if } u = e_2 \end{cases}.$$

Note that the ‘exit set’ Γ (see the Remark after Proposition 3) does not include (x, e_1) or $(x, -e_2)$, because these correspond to entering \mathcal{O} from its exterior. It may be checked that Q_b satisfies (7) on Γ . In order to make Q satisfy (7) and (4) on $\partial\mathcal{O} \times \mathcal{V}$, we have to set

$$Q_b(x, u, dv) = \begin{cases} \delta_{-e_1}(dv) & \text{if } u = -e_2, \\ \delta_{e_2}(dv) & \text{if } u = e_1 \end{cases}.$$

Conditions (3) and (5) are trivially satisfied on \mathcal{O} since $\lambda \equiv 0$.

The above is only a simple example of event chains in statistical mechanics rephrased as PDMPs; a systematic study is beyond the scope of this paper.

4 Subsampling

When using PDMC to sample from a posterior, we can use sub-samples of data at each iteration of the algorithm, as described in [5]. In the following we will assume that we can write the posterior as

$$\pi(x) \propto \prod_{i=1}^N f(y_i; x),$$

for some function f . For example this would be the likelihood for a single IID data point times the $1/N$ th power of the prior.

We first present a way for estimating this quantity in an unbiased manner, which uses control variates [2, 5]. For any $\hat{x} \in \mathcal{O}$ we note that

$$\nabla U(x) = \nabla U(\hat{x}) + [\nabla U(x) - \nabla U(\hat{x})].$$

We can then introduce the estimator $\widehat{\nabla U}(x)$ of ∇U by

$$\widehat{\nabla U}(x) = \nabla U(\hat{x}) + N [\nabla \log f(y_I; x) - \nabla \log f(y_I; \hat{x})], \quad (8)$$

where I is drawn uniformly from $\{1, \dots, N\}$.

The idea of using sub-sampling, within say the Bouncy Particle Sampler (BPS), is that at each iteration of our PDMC algorithm we can replace $\nabla U(x)$ by an unbiased estimator in step (3). We need to use the same estimate both when calculating the actual event rate in the accept/reject step and, if we accept, when simulating the new velocity. Whilst for each implementation of step (3) we need to use an independent estimate. The only further alteration we need to the algorithm is to choose an upper bound $\bar{\lambda}$ that holds

for all realizations of $\widehat{\nabla U}$. A more comprehensive explanation of this argument can be found in [12] in the context of the Zig-Zag sampler, and in [20] for the bouncy particle sampler.

It is straightforward to show that the resulting BPS algorithm uses an event rate that is $\mathbb{E} \left[\max \left(0, \widehat{\nabla U}(x) \cdot v \right) \right]$, and that this rate and the resulting transition probability Q at events satisfies Proposition 2 and 7. Hence this algorithm still targets $\pi(x)$, but only requires access to one data point at each accept-reject decision.

Note that this gain in computational efficiency does not come for free, as it follows from Jensen’s inequality that the overall rate of events will be higher. This makes mixing of the PDMC process slower. It is also immediate that the bound, $\bar{\lambda}$, we will have to use will be higher. However [5] show that if our estimator of $\widehat{\nabla U}(x)$ has sufficiently small variance, then we can still gain substantially in terms of efficiency. In particular they give an example where the CPU cost effective sample size does not grow with N – by comparison all standard MCMC algorithms would have a cost that is at least linear in N .

To obtain such a low-variance estimator requires a good choice of \hat{x} , so that with high probability x will be closer to \hat{x} . This involves a preprocessing step to find a value \hat{x} close to the posterior mode, a preprocessing step to then calculate $\nabla U(\hat{x})$ is also needed.

We now illustrate how to find an upper bound on the event rate. Following [5], if we assume L is a uniform (in space and i) upper bound on the largest eigenvalue of the Hessian of U^i , and if $\|v\| = 1$

$$\begin{aligned} & \max \left(0, \left(\nabla U(\hat{x}) + N(\nabla U^i(x_t) - \nabla U^i(\hat{x})) \right) \cdot v \right) \\ & \leq \max(0, \nabla U(\hat{x}) \cdot v) + N \left\| \nabla U^i(x) - \nabla U^i(\hat{x}) \right\| \\ & \quad + N \left\| \nabla U^i(x) - \nabla U^i(x_t) \right\| \\ & \leq \max(0, \nabla U(\hat{x}) \cdot v) + NL \|x - \hat{x}\| + NLt \end{aligned} \quad (9)$$

Thus the upper bound on the intensity is of the form $\bar{\lambda}(\tau) = a + b \cdot \tau$ with $a, b \geq 0$. In this case the first arrival time can be simulated as follows

$$\tau' = -a/b + \sqrt{\left(\frac{a}{b}\right)^2 + 2 \cdot \frac{R}{b}} \text{ with } R \sim \text{Exp}(1).$$

5 Experiments

We use Bayesian binary logistic regression as a testbed for our newly proposed methodology and perform a simulation study. The data $y_i \in \{-1, 1\}$ is modelled by

$$p(y_i | \xi_i, \beta) = f(y_i \beta^T \xi_i) \quad (11)$$

where $\xi \in \mathbb{R}^{p \times n}$ are fixed covariates and $f(z) = \frac{1}{1 + \exp(-z)} \in [0, 1]$. We will assume that we wish to fit this model under some monotonicity constraints – so that the probability of $y = 1$ is known to either increase or decrease with certain covariates. This is modeled through the constraint $\beta_i > 0$ and $\beta_i < 0$ respectively. This can be motivated for general logistic regression for questionnaires, see [25]. In following we consider the case $\beta_j \geq 0$ for $j = 1, \dots, p$.

For simplicity we use a flat prior over the space of parameters values consistent with our constraints. By Bayes’ rule the posterior π satisfies

$$\pi(\beta) \propto \prod_{i=1}^N f(y_i \beta^T \xi_i) \text{ for } \beta \in \mathcal{O},$$

where \mathcal{O} is the space of parameter values consistent with our constraints. We implement the BPS with subsampling, see Section 4 and use reflection at the boundary i.e. $Q_b(s, v, du) = \delta_{(I - 2P_{n(s)})v}(du)$ for $s \in \partial\mathcal{O}$. We can bound the switching intensity, even when we use the sub-sampling estimator of Equation (8), by a linear function of time. This follows because we can bound the second derivative of the log-likelihood of any observation. The details for deriving an upper bound on the intensity can be found in Section 7.2 of the supplementary material. We use $n = 10000$ and $p = 20$ and generate artificial data based on β^* and ξ whose components are a realization i.i.d. uniformly distributed random variables on $[0, 1]$. In Figure 3 we compare standard HMC with BPS in terms of effective sample size (ESS) per gradient evaluation. For HMC we use 5 leap-frog steps and tune the acceptance rate to 0.6. This requires a small step size because if HMC proposes outside the boundary the move is rejected. (There exists a version of HMC which can sample from truncated Gaussian distributions [21] but to our knowledge no efficient HMC scheme is available for general restricted domains.) We note that due to increasingly many rejections at the boundary, the effective sample size of the HMC scheme is relatively unchanged as the number of gradient evaluations is increased. On the other hand, we observe that the BPS shows a clear advantage over the HMC scheme in this case, with the effective sample size growing commensurately with the number of gradient evaluations.

6 Discussion

PDMC methods are a novel and promising alternative to existing classical discrete time MCMC methods. Due to their inherent non-reversibility, PDMC methods enjoy superior performance to their reversible




Figure 3: ESS against gradient evaluations for BPS and HMC (acceptance rate is hand tuned to 0.6) applied to logistic regression with $p = 20$ and $n = 10000$ and parameter β constrained to be positive. The graphic is based on 40 independent runs for each HMC and BPS for each choice of number of gradient evaluations.

counterparts. Other benefits, such as their amenability to unbiased subsampling techniques and their ability to leverage structural properties of the target distribution (such as any factor graph representation) makes them an extremely powerful tool in the Bayesian context, particularly in the large data regime.

In this work we have shown that PDMC methods also provide a very effective tool for sampling from probability distributions on restricted domains, which arise naturally in the Bayesian context where prior knowledge imposes hard constraints on inferred parameters, examples of which arise in econometrics and image processing.

While this work provides a framework for describing a general class of PDMC methods which are ergodic with respect to a given target probability distribution, it leaves open the question of how the choice of intensity function, velocity transition kernel as well as other parameters of the system influence the overall performance of the scheme. The problem of understanding the true computational cost of such PDMC schemes is more subtle than for classical discrete time MCMC schemes. Indeed, a PDMC scheme which maximizes rate of convergence to equilibrium or the speed of mixing need not be an optimal choice if the average switching rate (which is proportional to the number

of gradient evaluations per unit time), is very high. Further investigation is required to understand this delicate balance.

Acknowledgements

The authors wish to thank Thibaut Lienart for reading an early version of this paper, and for useful discussion. All authors thank the Alan Turing Institute and Lloyds registry foundation for support. S.J.V. gratefully acknowledges funding through EPSRC EP/N000188/1. J.B., P.F. and G.R. gratefully acknowledge EPSRC ilike grant EP/K014463/1. A.B.D. acknowledges grant EP/L020564/1.

References

- [1] J. H. Albert and S. Chib. Bayesian analysis of binary and polychotomous response data. *Journal of the American statistical Association*, 88(422):669–679, 1993.
- [2] R. Bardenet, A. Doucet, and C. Holmes. On markov chain monte carlo methods for tall data. *arXiv preprint arXiv:1505.02827*, 2015.
- [3] S. Bellavia, M. Macconi, and B. Morini. An interior point newton-like method for non-negative least-squares problems with degenerate solution. *Numerical Linear Algebra with Applications*, 13(10):825–846, 2006.
- [4] E. P. Bernard, W. Krauth, and D. B. Wilson. Event-chain Monte Carlo algorithms for hard-sphere systems. *Physical Review E*, 80(5):56704, nov 2009.
- [5] J. Bierkens, P. Fearnhead, and G. Roberts. The Zig-Zag Process and Super-Efficient Sampling for Bayesian Analysis of Big Data, 2016.
- [6] J. Bierkens and G. Roberts. A piecewise deterministic scaling limit of lifted metropolis-hastings in the curie-weiss model. *arXiv preprint arXiv:1509.00302*, 2015.
- [7] A. Bouchard-Côté, S. J. Vollmer, and A. Doucet. The Bouncy Particle Sampler: A Non-Reversible Rejection-Free Markov Chain Monte Carlo Method. *arXiv preprint arXiv:1510.02451*, 2015.
- [8] M. H. A. Davis. Piecewise-deterministic markov processes: A general class of non-diffusion stochastic models. *Journal of the Royal Statistical Society. Series B (Methodological)*, pages 353–388, 1984.
- [9] P. Diaconis, S. Holmes, and R. Neal. Analysis of a nonreversible Markov chain sampler. *Annals of Applied Probability*, 10(3):726–752, 2000.
- [10] S. Duane, A.D. Kennedy, B. J. Pendleton, and D. Roweth. Hybrid Monte Carlo. *Physics Letters B*, 195(2):216–222, 1987.
- [11] L. Fahrmeir and G. Tutz. *Multivariate statistical modelling based on generalized linear models. 2001*. Springer, New York.
- [12] P. Fearnhead, J. Bierkens, M. Pollock, and G. O. Roberts. Piecewise deterministic markov processes for continuous-time monte carlo. *arXiv preprint arXiv:1611.07873*, 2016.
- [13] N. Galbraith. On Event-Chain Monte Carlo Methods. Master’s thesis, Department of Statistics, Oxford University, 2016.
- [14] J. Geweke. Exact inference in the inequality constrained normal linear regression model. *Journal of Applied Econometrics*, 1(2):127–141, 1986.
- [15] Y. Guo and M. Berman. A comparison between subset selection and l1 regularisation with an application in spectroscopy. *Chemometrics and Intelligent Laboratory Systems*, 118:127–138, 2012.
- [16] J. Jacod and A. Shiryaev. *Limit theorems for stochastic processes*, volume 288. Springer Science & Business Media, 2013.
- [17] H. Kim and H. Park. Sparse non-negative matrix factorizations via alternating non-negativity-constrained least squares for microarray data analysis. *Bioinformatics*, 23(12):1495–1502, 2007.
- [18] P. A. Lewis and G. S. Shedler. Simulation of nonhomogeneous poisson processes by thinning. *Naval Research Logistics Quarterly*, 26(3):403–413, 1979.
- [19] Y. Nishikawa, M. Michel, W. Krauth, and K. Hukushima. Event-chain algorithm for the Heisenberg model. *Physical Review E*, 92(6):63306, 2015.
- [20] A. Pakman, D. Gilboa, D. Carlson, and L. Paninski. Stochastic bouncy particle sampler. *arXiv preprint arXiv:1609.00770*, 2016.
- [21] A. Pakman and L. Paninski. Exact hamiltonian monte carlo for truncated multivariate gaussians. *Journal of Computational and Graphical Statistics*, 23(2):518–542, 2014.

- [22] S. Patterson and Y. W. Teh. Stochastic gradient riemannian langevin dynamics on the probability simplex. In *Advances in Neural Information Processing Systems*, pages 3102–3110, 2013.
- [23] E. A. J. F. Peters and G. De With. Rejection-free Monte Carlo sampling for general potentials. *Physical Review E*, 85(2):1–5, 2012.
- [24] L. E. Train. *Discrete choice methods with simulation*. Cambridge university press, 2009.
- [25] Gerhard Tutz and Jan Gertheiss. Rating scales as predictors—the old question of scale level and some answers. *Psychometrika*, 79(3):357–376, 2014.
- [26] M. Welling and Y. W. Teh. Bayesian learning via stochastic gradient Langevin dynamics. In *Proceedings of the 28th International Conference on Machine Learning (ICML-11)*, pages 681–688, 2011.

7 Supplementary Material

In the supplementary material we shall provide a theoretical background for the framework described in Section 2. We first establish that the Markov process Z_t exists and is finite for all positive times.

Proof of Proposition 1. Suppose that \mathcal{V} is a compact subset of \mathbb{R}^d , we show that Z_t is non-explosive, which is equivalent to the fact that $\mathbb{E}N(t) < \infty$ for all $t > 0$, where $N(t)$ is the number of switching events that Z_t undergoes in the time $[0, t]$. Since \mathcal{V} is compact, this implies there exists $r_{max} > 0$ such that $\|v\| \leq r_{max}$, for all $v \in \mathcal{V}$. Therefore, we have $x_s \in B(x_0, tr_{max})$, for all $s \in [0, t]$. Since $\sum_{i=1}^N \lambda_i(\cdot, v)$ is piecewise continuous for all $v \in \mathcal{V}$, this implies that there exists $\lambda_{max} > 0$ such that $\sum \lambda_i(x_s, v_s) < \lambda_{max}$, for $s \in [0, t]$. The expected number of switches up to time t can thus be written as

$$\mathbb{E}N(t) = \int_0^t \sum_{i=1}^N \lambda_i(X_s, V_s) ds \leq \lambda_{max} t < \infty,$$

giving the required result. \square

7.1 Sampling from domains without restrictions

Let us first focus on the case where $\mathcal{O} = \mathbb{R}^d$. From [8, Section 5] the process Z_t will have infinitesimal generator given by the closure of the operator

$$\mathcal{L}f(x, v) = v \cdot \nabla_x f(x, v) + \sum_{i=1}^N \lambda_i(x, v) \int_{\mathcal{V}} (f(x, u) - f(x, v)) Q_i(x, v, du), \quad (x, v) \in E, \quad (12)$$

where $\mathcal{D}(\mathcal{L})$ is the set of functions which are continuously differentiable with respect to x and which decay sufficiently fast as $\|x\| \rightarrow \infty$. Having characterized the infinitesimal generator, we can now prove Proposition 2.

Sketch Proof of Proposition 2. Without loss of generality we take $\pi(x) = \exp(-U(x))$, i.e. the proportionality factor in $\pi(x) \propto \exp(-U(x))$ is assumed to be 1. We shall only provide a formal proof of this result, by demonstrating that

$$\int \mathcal{L}f(x, v) \pi(x) dx \rho(dv) = 0, \quad \text{for all } f \in \mathcal{D}(\mathcal{L}),$$

so that \mathcal{L} is infinitesimally invariant. A rigorous proof would require establishing that $\mathcal{D}(\mathcal{L})$ as defined above is a core for the extended generator. This is a technical result which we defer for future work. Let $f \in \mathcal{C}$, then we have that

$$\begin{aligned} & \int_{\mathbb{R}^d} \int_{\mathcal{V}} \int_{\mathcal{V}} \sum_{i=1}^N \lambda_i(x, v) [f(x, u) - f(x, v)] Q_i(x, v, du) \pi(x) dx \rho(dv) \\ &= \int_{\mathbb{R}^d} \int_{\mathcal{V}} \int_{\mathcal{V}} \sum_{i=1}^N \lambda_i(x, v) f(x, u) Q_i(x, v, du) \rho(dv) \pi(x) dx - \int_{\mathbb{R}^d} \int_{\mathcal{V}} \int_{\mathcal{V}} \sum_{i=1}^N \lambda_i(x, v) f(x, v) Q_i(x, v, du) \rho(dv) \pi(x) dx \\ &= \int_{\mathbb{R}^d} \int_{\mathcal{V}} \int_{\mathcal{V}} \sum_{i=1}^N \lambda_i(x, v) f(x, u) Q_i(x, v, du) \rho(dv) \pi(x) dx - \int_{\mathbb{R}^d} \int_{\mathcal{V}} \sum_{i=1}^N \lambda_i(x, v) f(x, v) \rho(dv) \pi(x) dx \\ &= \int_{\mathbb{R}^d} \int_{\mathcal{V}} \sum_{i=1}^N f(x, u) \left[\int_{\mathcal{V}} \lambda_i(x, v) Q_i(x, v, dv) - \lambda_i(x, u) \right] \rho(du) \pi(x) dx \\ &= - \int_{\mathbb{R}^d} \int_{\mathcal{V}} f(x, u) u \cdot \nabla U(x) \rho(du) \pi(x) dx \\ &= \int_{\mathbb{R}^d} \int_{\mathcal{V}} f(x, u) u \cdot \nabla \pi(x) \rho(du) dx \\ &= - \int_{\mathbb{R}^d} \int_{\mathcal{V}} u \cdot \nabla_x f(x, u) \pi(x) \rho(du) dx. \end{aligned}$$

It follows that

$$\int_{\mathbb{R}^d} \int_{\mathcal{V}} \mathcal{L}f(x, v) \pi(x) dx \rho(dv) = \int_{\mathbb{R}^d} \int_{\mathcal{V}} (u \cdot \nabla_x f(x, u) - u \cdot \nabla_x f(x, u)) \pi(x) dx \rho(dv) = 0,$$

so that $\pi(x) dx \otimes \rho(dv)$ is infinitesimally invariant with respect to Z_t . \square

7.1.1 Sampling from domains with restrictions

We now consider the case where \mathcal{O} is a pathwise connected, open subset of \mathbb{R}^d with Lipschitz boundary $\partial\mathcal{O}$. Following [8], the PDMP Z_t will have infinitesimal generator given by the closure of (12) with domain $\mathcal{D}(\mathcal{L})$ of continuously differentiable functions vanishing sufficiently fast as $\|x\| \rightarrow \infty$, such that

$$f(x, v) = \int_{\mathcal{V}} f(x, u) Q_b(x, v, du), \quad (13)$$

for all $(x, v) \in \partial\mathcal{O} \times \mathcal{V}$. Based on this identification of the infinitesimal generator we can now provide a formal proof that the conditions of Proposition 3 are sufficient to ensure invariant of $\pi \otimes \rho$.

Sketch Proof of Proposition 3. Following the argument of the proof of Proposition 1, for $f \in \mathcal{D}$, we can write

$$\begin{aligned} & \int_{\mathcal{O}} \int_{\mathcal{V}} \int_{\mathcal{V}} \sum_{i=1}^N \lambda_i(x, v) [f(x, u) - f(x, v)] Q_i(x, v, du) \rho(dv) \pi(x) dx \\ &= - \int_{\mathcal{O}} \int_{\mathcal{V}} u \cdot \nabla f(x, u) e^{-U(x)} \rho(du) dx + \int_{\partial\mathcal{O}} \int_{\mathcal{V}} f(\sigma, u) (u \cdot n(\sigma)) e^{-U(\sigma)} \rho(du) d\sigma, \end{aligned}$$

where the boundary term arises from integration by parts with respect to x . Considering the boundary integral, by applying (4) which is assumed to hold on $\partial\mathcal{O}$ and (13) we obtain

$$\begin{aligned} & \int_{\partial\mathcal{O}} \int_{\mathcal{V}} f(\sigma, u) (u \cdot n(\sigma)) e^{-U(\sigma)} \rho(du) d\sigma \\ &= \int_{\partial\mathcal{O}} \int_{\mathcal{V}} \int_{\mathcal{V}} f(\sigma, v) Q_b(\sigma, u, dv) (u \cdot n(\sigma)) e^{-U(\sigma)} \rho(du) d\sigma \\ &= \int_{\partial\mathcal{O}} \int_{\mathcal{V}} \int_{\mathcal{V}} f(\sigma, v) Q_b(\sigma, v, du) (u \cdot n(\sigma)) e^{-U(\sigma)} \rho(dv) d\sigma \\ &= - \int_{\partial\mathcal{O}} \int_{\mathcal{V}} f(\sigma, v) (v \cdot n(\sigma)) e^{-U(\sigma)} \rho(dv) d\sigma, \end{aligned}$$

so that the boundary term evaluates to zero. The argument then continues as in Proposition 2, thus establishing the invariance of the target distribution. \square

Sketch Proof of Proposition 4. Let $f \in \mathcal{D}(\mathcal{L})$, so that f satisfies (13). By the assumptions on Q_b in Proposition 4, this implies that $f(x, v) = f(x)$ for all $x \in \partial\mathcal{O}$. Following the proof of Proposition 3, the boundary integral term becomes

$$\int_{\partial\mathcal{O}} \int_{\mathcal{V}} f(s, u) (u \cdot n(s)) e^{-U(s)} \rho(du) ds = \int_{\partial\mathcal{O}} f(s) e^{-U(s)} n(s) ds \cdot \int_{\mathcal{V}} u \rho(du),$$

which is zero if ρ has mean zero, as required. \square

7.2 Derivation of dominating intensity for logistic regression example

A valid choice of L can be derived as follows: Notice that $(\log f(z))' = f(-z)$ and $f'(z) = f(z)(1 - f(z))$ so that we obtain

$$\begin{aligned} \frac{\partial}{\partial x} \log f(y_i | x) &= f(-y_i x^\top \xi_i) y_i \xi_i \\ \frac{\partial}{\partial^2 x} \log f(y_i | x) &= -f(-y_i x^\top \xi_i) (1 - f(-y_i x^\top \xi_i)) \xi_i \xi_i^\top \end{aligned}$$

Using $p(1-p) \leq \frac{1}{4}$ for $p \in [0, 1]$

$$\sup_{\|w\| \leq 1} \left| w^t \frac{\partial}{\partial^2 x} \log f(y_i|x) w^t \right| \leq \frac{1}{4} \|\xi_i\|^2$$

Such that Equation 10 holds with

$$L := \frac{1}{4} \max_{i=1, \dots, n} \|\xi_i\|$$

as defined above.